

Introduction: Research in planetary seismology is fundamentally constrained by a lack of data. Seismological science products of future missions can typically only be informed by theoretical signal/noise characteristics of the environment [1] or likely Earth-analogues [2]. Although objectives can be re-assessed after some initial data-collection upon lander arrival, transfer of high-resolution data back to Earth is costly on lander power usage.

Over the last several years, development of GPU computing techniques and open-source high-level APIs have led to rapid advances in deep learning within the fields of computer vision, natural language processing, and collaborative filtering. These techniques are actively being adapted in seismology for a variety of tasks, including: earthquake detection [3], seismic phase discrimination [4], and ground-motion prediction [5].

Until the recent detection of marsquakes during the Mars InSight mission, the only other measurements of seismicity recorded outside of Earth was on the Moon during the Apollo missions between 1969 to 1977 [6]. These unique datasets have been periodically revisited using new seismological methods, including ambient noise interferometry [7] and Hidden Markov Models [8].

Our objective is to develop a deep learning seismic detector and use it to catalog moonquakes from the Apollo 17 Lunar Seismic Profiling Experiment (LSPE) and compare the results with those obtained by other methods. Additionally, we will assess the accuracy tradeoff between using a training set of lunar data and one composed of Earth seismicity. In this document, we present preliminary results using a prototype classifier trained on a small set of earthquakes that was able to obtain detections for LSPE moonquakes with a greater accuracy than a recent study using Hidden Markov Models [9].

Methodology: We built a prototype deep learning classifier that was able to distinguish between seismic activity and noise through examples of spectrogram images for each category recorded on Earth. This process will be expanded to two classifiers, one using a more comprehensive arrangement of Earth seismic data and the other using moonquakes and lunar seismic noise.

Earth seismic data was downloaded using the *IRIS* utility *PyWeed* in a time window around the earthquake first arrival. For the prototype classifier, we used an interval of 180 seconds before and 20 seconds after the

P-wave arrival of earthquakes greater than M_w 3 from the Piñon Flats Observatory (PFO) seismic station [Fig. 1].

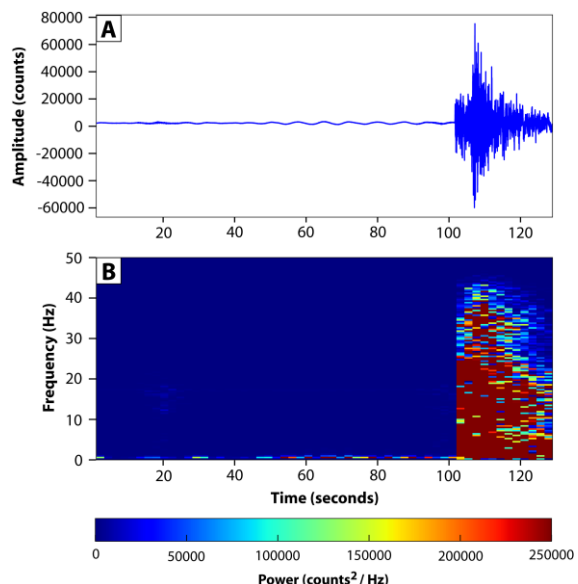


Figure 1: Time series [A] and spectrogram [B] of a M_w 3.7 earthquake at PFO.

Data augmentation is a technique in computer vision where new images are created by modifying existing data [10]. In image recognition, this is typically done by cropping, zooming, or rotating images. However, applying data augmentation in this manner will decrease the accuracy of our model, as cropped spectrograms may omit valuable information in the low or high frequencies. Instead, we chose twenty sliding windows across the noise and earthquake segments with one second overlap starting at 0 seconds for the noise and at 81 seconds for the earthquake (19 seconds prior to the onset of the P-wave at 101 seconds). A total of 27,800 spectrograms were used in the prototype, approximately 20% of which (5240) were separated into a validation set.

The prototype was built using the *fastai* computer vision library [10] with a batch size of 16 and image reduction to 224x224 pixels. Three training cycles were conducted on the data with a learning rate of $1e-2$, which took approximately 18 minutes on a laptop with an NVIDIA Quadro M1200 GPU (4 Gigabytes of video memory). We used the ResNet 34 CNN architecture [11] for this prototype, but we will determine an

optimal and unique architecture using best practices from literature for the full experiment.




Figure 2: Time series [A] and spectrogram [B] of a moonquake recorded on February 11th 1977 during the Apollo 17 LPSE.

Preliminary Results: The prototype model obtained 94.5% accuracy on the 5240 element validation set after three rounds of training (training duration was approximately 6 minutes per round). The loss for the training and validation sets is the summation of the errors in the cross-entropy, also known as the negative-log likelihood [10]. Additional rounds only reduced the training loss but not the validation loss, which suggests that further training overfits the data.

This model was applied to a 2.5 hour data segment collected by the LSPE on February 11th. Two pre-processing steps were necessary to account for differences between the Earth-trained model and the lunar data: (1) the moonquake signal is much weaker in power compared to earthquakes and had to be capped at $1e-6$ counts²/Hz, and (2) a 10 Hz highpass filter was applied to remove low-frequency lunar noise. The classifier found a total of five detections after a running time of approximately 20 minutes. The obtained detections are generally of equal quality or more accurate than the results obtained using Hidden Markov Models and do not contain any false detection. However, the classifier appears to have difficulty detecting the shorter duration signals that occur in the last hour of the record [Figure 3]. We will build a new model using a unique architecture and comprehensively quantify the number of moonquakes in the Apollo 17 LSPE dataset. The same experiment will be repeated using a new model trained from lunar seismic data, and we will

assess the accuracy tradeoff from using non-local training data.

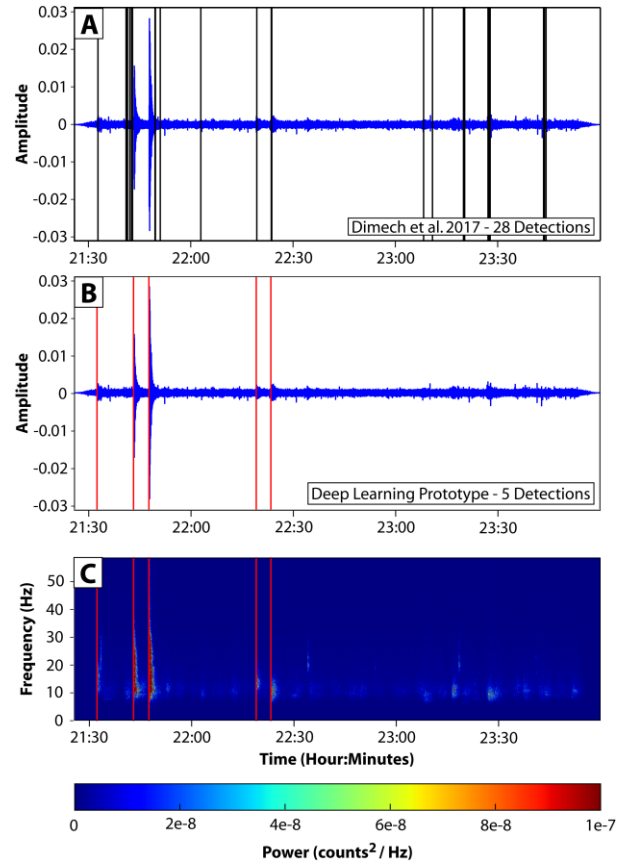


Figure 3: Detection comparison between the Hidden Markov Model study [9] [A, black lines] and the deep learning prototype developed for this proposal [B, C, red lines].

References:

- [1] Panning M. et al. (2018) *JGR: Planets*, 123, 167-179.
- [2] Zhan, Z. et al. (2014) *GJI*, 196, 1796-1802.
- [3] Meier, M. et al. (2019) *JGR Solid Earth*, 124, 788-800.
- [4] Ross, Z. et al. (2019) *JGR Solid Earth*, 124, 856-869.
- [5] Trugman, D. and Shearer, P. (2018) *BSSA*, 108, 929-945.
- [6] Lognonné, P. (2005) *Annual Reviews of Earth Planet Science*, 33, 571-604.
- [7] Larose et al. (2005) *GRL*, 32 (16).
- [8] Knapmeyer-Endrun, B. and Hammer, C. (2015) *JGR: Planets*, 120, 1620-1645.
- [9] Dimech, J. et al. (2017) *Results in Physics*, 7, 4457-4458.
- [10] Howard, J. et al. (2018) *fastai*.
- [11] He, K. et al. (2016) *2016 IEEE Conference on Computer Vision and Pattern Recognition*.